

Fast Approximate Bayesian Computation for Estimating Parameters in Differential Equations

Sanmitra Ghosh, Srinandan Dasmahapatra, and Koushik Maharatna, *Member, IEEE*

Abstract—Approximate Bayesian computation (ABC) using a sequential Monte Carlo method provides a comprehensive platform for parameter estimation, model selection and sensitivity analysis in differential equations. However, this method, like other Monte Carlo methods, incurs a significant computational cost as it requires explicit numerical integration of differential equations to carry out inference. In this paper we propose a novel method for circumventing the requirement of explicit integration by using derivatives of Gaussian processes to smooth the observations from which parameters are estimated. We evaluate our methods using synthetic data generated from model biological systems described by ordinary and delay differential equations. Upon comparing the performance of our method to existing ABC techniques, we demonstrate that it produces comparably reliable parameter estimates at a significantly reduced execution time.

Index Terms—Approximate Bayesian computation, Gaussian process regression, non-linear differential equations, non-parametric Bayesian, sequential Monte Carlo

1 INTRODUCTION

THE time evolution of the variables modelled in a variety of science and engineering branches are often described by ordinary differential equations that are characterised by model structure – the functions of the dynamical variables – and model parameters. The task of estimating these parameters from experimental observations is thus of paramount importance. It is also necessary in some cases to choose the most appropriate among competing models that describe the observations. For parameter estimation and model selection, statistical and pattern recognition techniques built upon a Bayesian framework have been shown to work extremely well for complex non-linear ordinary differential equations in [1],[2],[3] and [4].

To apply Bayesian techniques we need to integrate marginal likelihoods, which can be computationally intractable in non-linear differential equation models. For this reason some form of approximation such as Monte Carlo integration is generally preferred for parameter inference. Approximate Bayesian computation (ABC) based on sequential Monte Carlo (SMC) is one such approximate inference technique that has been applied to different classes of dynamical systems described by deterministic or stochastic differential equations for both parameter estimation and model selection in [5]. The ABC-SMC algorithm has been shown to work well for the examples considered in [5]. ABC-SMC produces reliable estimates of parameters and has been used to discriminate between a set of candidate models using Bayesian model selection criteria. Moreover, ABC-SMC

enables the calculation of parameter sensitivities [5].

ABC methods prove to be most useful for large models with complex likelihood surfaces that are difficult to evaluate. The operating principle of ABC methods lies in replacing the evaluation of likelihoods with a simulation based procedure for inference, by using a generative model \mathcal{M}_θ with parameters θ drawn from a prior distribution $\pi(\theta)$ to simulate observations $Y^s \sim \mathcal{M}_\theta$ that are compared with the observed data Y^d . If the likelihood $p(Y^d|\mathcal{M}_\theta)$ of observed data Y^d is intractable or infeasible to compute, then we can use the ABC algorithm to obtain samples from the following modified posterior density

$$p_\epsilon(\theta, Y^s|Y^d) = \frac{\mathbf{1}(\Delta(Y^d, Y^s) \leq \epsilon)(Y^s \sim \mathcal{M}_\theta)\pi(\theta)}{\int_\theta \int_{Y^s} \mathbf{1}(\Delta(Y^d, Y^s) \leq \epsilon)(Y^s \sim \mathcal{M}_\theta)\pi(\theta)d\theta dY^s} \quad (1)$$

where $\epsilon > 0$ is a tolerance level, Δ is a distance function, $\mathbf{1}$ is the indicator function and $p_\epsilon(\theta, Y^s|Y^d) = p(\theta, Y^s|\Delta(Y^d, Y^s) \leq \epsilon)$. A good (enough) approximation of the true marginal posterior distribution is obtained when the distance $\Delta(Y^d, Y^s)$ is within a predetermined small tolerance ϵ , i.e.,

$$p_\epsilon(\theta|Y^d) = \int_{Y^s} p_\epsilon(\theta, Y^s|Y^d)dY^s \approx p(\theta|Y^d) \quad (2)$$

Since ABC (including ABC-SMC) requires the generation of a number of simulated observations $Y^s \sim \mathcal{M}_\theta$, the generation of observations could be a computationally expensive process. Thus although ABC-SMC mitigates the intractability of evaluating the likelihood function through simulation, repeated simulation from complex models for inference can itself be burdensome. For the case of dynamical systems such simulations require explicit numerical solutions of non-linear differential equations. Thus, despite its attractive features the ABC-SMC algorithm suffers from a major drawback rooted

• S. Ghosh, S. Dasmahapatra and K. Maharatna are with the Department of Electronics and Computer Science, University of Southampton, Southampton, SO17 1BJ.
E-mail: {sg5g10, sd, km3}@ecs.soton.ac.uk

in its computational burden for inference in differential equations. In particular, the acceptance criterion $\mathbb{1}(\Delta(Y^d, Y^s) \leq \epsilon)$ can lead to the generation of many unused simulations $Y^s \sim \mathcal{M}_\theta$, and various methods have been proposed in [6] to improve the acceptance rate and reduce the run-time of the algorithm.

In this paper we propose an alternate method of speeding up the ABC-SMC algorithm for parameter estimation in deterministic models described by ordinary differential equations (ODE) or delay differential equations (DDE) by reducing the time incurred in simulation. We achieve this speedup by: (i) completely circumventing the process of integrating the differential equation by operating on the derivative space and (ii) by smoothing the derivatives using Gaussian processes (GP). It should be noted that using Gaussian processes as functional emulators in the derivative space, as a concept, has been proposed in [1],[2] for speeding up parameter estimation in deterministic differential equations. For parameter estimation, GP-based gradient matching has been used for ODEs and DDEs using a population Monte Carlo sampling [1]; an adaptive variant of this approach is proposed for ODEs [2]. See [3] for a review and comparison between these approaches. The novelty of our proposed method is the fusion of GP regression with ABC-SMC. Our algorithm for fast parameter estimation can be easily incorporated into methods for model selection and recovering parameter sensitivities for deterministic differential equations.

This paper is organised as follows: in section 2 we will introduce the ABC-SMC algorithm for parameter estimation in differential equations. In section 3 we will show how the algorithm can be sped up by circumventing the need for the numerical solution of the differential equations. In section 4 we will introduce Gaussian processes for function estimation and subsequently in section 5 we will use GPs within the ABC-SMC to operate on the derivative space. In section 6 we will compare the performance of our proposed modification with the ABC-SMC algorithm [5] and also with its improved variant proposed in [6]. In conclusion, we summarise our achievements in this paper and give some indication of future work in section 7.

2 THE BASIC ABC FRAMEWORK

ABC methods generally have the following algorithmic form:

- 1) Sample a candidate parameter vector θ from a prior distribution $\pi(\theta)$ and for each $\theta \sim \pi(\theta)$, simulate a dataset $Y^s \sim \mathcal{M}_\theta$ from a generative model \mathcal{M}_θ .
- 2) Compute a distance $\Delta(Y^s, Y^d)$ between the simulated dataset, Y^s and the experimental data Y^d . If $\Delta(Y^d, Y^s) \leq \epsilon$, where $\epsilon \geq 0$ is the error tolerance of accepted solutions, then accept θ and reject otherwise.

These scheme is repeated until N parameter values are accepted, which represent a sample from the approxi-

mate posterior distribution $p_\epsilon(\theta|Y^d)$. Exact posterior can be obtained from this scheme when $\epsilon = 0$.

2.1 ABC-SMC for parameter estimation in ODE

If the prior distribution is very different from the posterior distribution, the basic ABC framework is very inefficient as it spends a considerable amount of time sampling from areas of low likelihood in parameter space, which makes the acceptance rate extremely low. In order to improve upon poor acceptance rates and facilitate exploration of the parameter space, ABC algorithms based on the SMC sampling method were proposed in [7], [8] and [9] and sequential importance sampling (SIS) in [5], [10]. [5] applied the ABC algorithm based on SIS for parameter estimation and model selection for a variety of dynamical systems including non-linear ODEs and DDEs, which will also be the focus in this paper. Although all the variants of ABC algorithms that come under the SMC category can potentially be used for inference in dynamical systems, we will specifically focus our attention on the ABC approach as adopted in [5].

We shall apply ABC-SMC to models of the evolution of state $\mathbf{X}(t) = (X_1(t), \dots, X_K(t))$ that are governed by ODEs or DDEs $\frac{d\mathbf{X}(t)}{dt} = f(\mathbf{X}(t - t_d), \theta)$, where t_d stands for the time delay in DDEs, with $t_d = 0$ for ODEs, and θ is a vector of parameter values. We express the integrated solution of the differential equations $\mathbf{X}(t, \mathbf{X}_{\text{in}}; \theta)$ as a map $\psi_t(\mathbf{X}_{\text{in}}; \theta)$ that generates state trajectories $\mathbf{X}(t)$ given a set of parameters θ and initial conditions $\mathbf{X}_{\text{in}} \triangleq \mathbf{X}(t_d \leq t \leq 0)$. To generate the samples Y^s , we add Gaussian noise $\eta \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbb{I}_K)$, where \mathbb{I}_K is a $K \times K$ identity matrix to the solutions $\mathbf{X}(t, \mathbf{X}_{\text{in}}; \theta)$ to create the generative model \mathcal{M}_θ :

$$(Y^s \sim \mathcal{M}_\theta) \Leftrightarrow Y^s = \mathbf{X}(t, \mathbf{X}_{\text{in}}; \theta). \quad (3)$$

to be used in the ABC framework.

A collection of parameter values, called particles θ are sampled from the prior $\pi(\theta)$ to instantiate the generative model \mathcal{M}_θ . To decide whether a particular choice of parameters $\theta \sim \pi(\theta)$ is accepted, we need to compare if the simulated trajectory $Y^s \sim \mathcal{M}_\theta$ is within a tolerance level ϵ of the observed trajectory Y^d , for which we introduce the distance function

$$\Delta(Y^d, Y^s) = \sum_{i=1}^L \sum_{k=1}^K (Y_k^d(t_i) - X_k(t_i) - \eta_k)^2, \quad (4)$$

where we assumed that the data were collected, and the state evaluated, at integer time points $t^L \triangleq \{t_i\}_{i=1, \dots, L}$. Note that for dynamical systems such distance functions are generally built by considering the entire time-series data instead of some sufficient statistics.

The sequential stage of this algorithm involves replacing a single tolerance value ϵ by a sequence of tolerance values ϵ_τ , where $\tau = 0, \dots, S_{MC}$ denotes the sequential steps, and $\epsilon_\tau > \epsilon_{\tau+1}$. The particles θ_τ are indexed by τ labelling the tolerance level, and are sampled from

the posterior distribution obtained from the previous sequential step, thus introducing a step-wise procedure for generating parameters from a sequence of more informative distributions, starting at $\tau = 0$ with the prior distribution $\pi(\theta)$. To accept or reject the sampled particles for sequence index τ , the generated trajectories from the model with parameters θ_τ must be closer to the observed data Y^d than those generated from the model with parameters $\theta_{\tau-1}$ in step $\tau - 1$. The generative mechanism for the particles θ_τ differs from that of $\theta_{\tau-1}$ in that they are sampled from the N particles $\{\theta_{\tau-1}^{(i)}\}_{i=1,\dots,N}$ with importance weights [5] $w_{\tau-1}$ and each $\theta^* \sim \{\theta_{\tau-1}^{(i)}\}_{i=1,\dots,N}$ is perturbed by a perturbation kernel $K_\tau(\theta|\theta^*)$ [5].

For each $\tau = 0, \dots, S_{MC}$, the N particles meeting the acceptance criterion $\Delta(Y^s, Y^d) \leq \epsilon_\tau$ represent a point-wise approximation for the posterior distribution over the parameter values:

$$p_{\epsilon_\tau}(\theta|Y^d) \approx \frac{1}{N} \sum_{i=1}^N w_\tau^{(i)} \delta(\theta - \theta_\tau^{(i)}), \quad (5)$$

where $w_\tau^{(i)}$ is the importance weight of the particle i [5] and $\delta(\theta - \theta^{(i)})$ is a product of Dirac delta functions, one for each component of θ . For a small value of $\epsilon_{S_{MC}}$ the final collection of particles should be a good point-wise approximation to the true posterior distribution. The ABC-SMC algorithm is listed in Algorithm 1.

One way of speeding up the ABC-SMC algorithm is by increasing the acceptance rate. To this end a range of perturbation kernels were proposed in [6] that result in a noticeable change in terms of acceptance rates and run-time. However, the major computational burden stems from the numerical integration of the differential equation and thus a faster simulation method within the ABC-SMC is likely to speed up this algorithm possibly more than what is gained by clever choice of perturbation kernels.

3 GRADIENT BASED PARAMETER ESTIMATION IN DIFFERENTIAL EQUATIONS

We have mentioned previously that the computational bottleneck stems from the explicit integration carried out in each simulation step. In order to avoid the integration one could essentially use a gradient based estimation. If the temporal variations in observations $Y^d(t)$ are believed to be less smooth than the underlying state evolution that is modelled by differential equations, we shall introduce the target state variable $\widehat{X}(t)$ to be the smoothed version $\widehat{X}(t) \triangleq \mathcal{S}(Y^d)$ of the observations. Here, \mathcal{S} represents any smoothing procedure, and we will use Gaussian Process (GP) regression to perform the smoothing below. In the ABC framework, we shall accept the trajectories $X(t)$ from the model \mathcal{M}_θ (see (3)) if they are close to $\widehat{X}(t)$. Once we have $\widehat{X}(t)$ we can compute its numerical derivative to obtain the empirical

Algorithm 1 ABC-SMC as proposed in [5]

1. Given Y^d , $\pi(\theta)$, \mathcal{M}_θ .
 2. Initialise $\epsilon_\tau > 0$, $\tau = 0, \dots, S_{MC}$, $\epsilon_\tau > \epsilon_{\tau+1}$. Set $\tau = 0$.
 3. Set $i = 1$.
 4. **if** $\tau = 0$ **then**
 5. sample θ^{**} independently from $\pi(\theta)$:
 $\theta^{**} \sim \pi(\theta)$
 6. **else**
 7. from the previous population $\{\theta_{\tau-1}^{(i)}\}_{i=1,\dots,N}$
 sample $\theta^* \sim \{\theta_{\tau-1}^{(i)}\}_{i=1,\dots,N}$ with associated normalized weights $w_{\tau-1}^*$ and use the perturbation kernel $K_\tau(\theta|\theta^*)$ to produce $\theta^{**} \sim K_\tau(\theta|\theta^*)$.
 8. **end if**
 9. **if** $\pi(\theta^{**}) = 0$ **then**
 10. go to 4.
 11. **else**
 12. Simulate a candidate dataset Y^s from the model \mathcal{M}_θ with parameter θ^{**} : $Y^s \sim \mathcal{M}_\theta|_{\theta \leftarrow \theta^{**}}$.
 13. **end if**
 14. **if** $\Delta(Y^d, Y^s) \geq \epsilon_\tau$ **then**
 15. go to 4.
 16. **else**
 17. Set $\theta_\tau^{(i)} \leftarrow \theta^{**}$ and calculate the weight for particle $\theta_\tau^{(i)}$,

$$w_\tau^{(i)} = \begin{cases} 1, & \text{if } \tau = 0 \\ \frac{\pi(\theta_\tau^{(i)})}{\sum_{j=1}^N w_{\tau-1}^{(j)} K_\tau(\theta_\tau^{(i)}|\theta_{\tau-1}^{(j)})}, & \text{if } \tau > 0 \end{cases}$$
 18. **end if**
 19. **if** $i < N$ **then**
 20. Set $i \leftarrow i + 1$ and go to 4.
 21. **else**
 22. Normalise the weights.
 23. **end if**
 24. **if** $\tau < S_{MC}$ **then**
 25. Set $\tau \leftarrow \tau + 1$ and go to 3.
 26. **else**
 27. **return** particles $\theta_{S_{MC}}^{(i)}$ at $\tau = S_{MC}$.
 28. **end if**
-

vector field $V^d(t)$ of the dynamical system \mathcal{M}_θ :

$$\text{for } \widehat{X}(t) \triangleq \mathcal{S}(Y^d), V^d(t) \triangleq \frac{d}{dt} \widehat{X}(t). \quad (6)$$

In addition, while the left hand side of the ODE $\frac{d}{dt} X(t) = f(X(t), \theta)$ is estimated by the empirical derivative $V^d(t)$, it should be matched by the model vector field $f(X(t), \theta)$ on the right hand side, when evaluated on the smoothed state data $f(X(t) = \widehat{X}(t), \theta)$.

Upon introducing a new distance measure between $V^d(t)$ obtained from the smoothing and $f(\widehat{X}(t), \theta)$ obtained from the vector field we can eliminate the original distance metric for ABC-SMC between observed, Y^d ,

and the simulated, $Y^s \sim \mathcal{M}_\theta$, trajectories, thus unbundling ABC-SMC of ODE integration at each simulation step. The gradient based method was first suggested in [11] where a spline-based smoothing was used to denoise the observed data. In this method a cost function was built using the distance metric in derivative space and optimisation was used to minimise this cost function in order to obtain point estimates. Recent developments of this methods are described in [12]. All these approaches suffer from similar problems of using additional regularisation parameters for smoothing and often the estimates are sub-optimal point estimates. Although porting the derivative based distance within an ABC scheme alleviates the computational bottleneck, this approach suffers from an inherent shortcoming that is rooted in obtaining a numerical derivative as this might lead to information loss.

In our approach, we replace the numerical differentiation with a zero mean Gaussian Process (GP) prior on the state $\mathbf{X}(t)$ given by

$$p(\mathbf{X}(t)|\phi) \sim \mathcal{GP}(0, K(t, t'; \phi)), \quad (7)$$

where $K(t, t'; \phi)$ denotes a covariance function with hyperparameters ϕ . Once such a prior is established then Gaussian Process regression techniques can be applied to estimate both the state vector $\widehat{\mathbf{X}}(t)$ and also the derivative process $\mathbf{V}^d(t)$. Using GP regression the derivative process can be inferred within a probabilistic framework. Hence, we propose to use a distance function in the derivative space where the state $\widehat{\mathbf{X}}(t)$ and derivative $\mathbf{V}^d(t)$ is modelled using GP regression, within the ABC-SMC algorithm. In this way our proposed method is based on the GP construction in the derivative space as in [1] and [3], combined with the ABC-SMC algorithm as proposed in [5]. Next, we will briefly introduce GP regression and will apply this to the ABC-SMC algorithm.

4 GAUSSIAN PROCESSES

For real-valued functions $f : \mathcal{A} \rightarrow \mathbb{R}$ of one or more input variables defined over \mathcal{A} , a Gaussian process (GP) is a Bayesian non-parametric model that specifies a distribution $P(f)$ [13], [14], [15], [16] characterised by a mean function $\mu(x)$ and a covariance function or kernel, $K(x, x'; \phi)$:

$$f(x) \sim \mathcal{GP}(\mu(x), K(x, x'; \phi)), \quad (8)$$

where ϕ are hyperparameters. For example, the squared exponential covariance function, which we use below, is given by

$$K_{SE}(x, x') = \sigma_{\text{kernel}}^2 \exp\left(-\frac{1}{2} \frac{(x - x')^2}{l^2}\right), \quad (9)$$

with hyperparameters σ_{kernel}^2 and l^2 (variance and characteristic length-scale).

For a finite number (n) of inputs $\mathbf{x}^* = (x_1, \dots, x_n)$, $x_i \in \mathcal{A}$ and for $f(x) \sim \mathcal{GP}(\mu(x), K(x, x'; \phi))$, the n -dimensional vector of function values evaluated at n

points $\mathbf{f}(\mathbf{x}^*) \triangleq (f(x_1), \dots, f(x_n))$ is a random vector drawn from a multivariate Gaussian distribution:

$$p(\mathbf{f}(\mathbf{x}^*)|\mathbf{x}^*) = \mathcal{N}(\mu(\mathbf{x}^*), K(\mathbf{x}^*, \mathbf{x}'^*)). \quad (10)$$

For performing regression, observations $\mathbf{y}(x) \triangleq (y(x_1), \dots, y(x_L))$ at L training points $\mathbf{x} = (x_1, \dots, x_L)$ are fit to function f evaluated at \mathbf{x} :

$$\mathbf{y}(x) = \mathbf{f}(x) + \boldsymbol{\eta}, \quad (11)$$

as in (3). Given the training data $(x_i, y(x_i))$, $i = 1, \dots, L$ the conditional predictive distribution of the function $\mathbf{f}^* \triangleq \mathbf{f}(\mathbf{x}^*)$ evaluated at the test points \mathbf{x}^* is a Gaussian with mean and variance given by [16]

$$\begin{aligned} E[\mathbf{f}^*|\mathbf{y}, \mathbf{x}, \mathbf{x}^*, \Phi] &= \mu(\mathbf{f}^*) + \\ &\quad [K(\mathbf{x}^*, \mathbf{x})(K(\mathbf{x}, \mathbf{x}) + \sigma^2 \mathbb{I})^{-1} \\ &\quad \times (\mathbf{y} - \mu(\mathbf{f}^*))], \\ \text{Var}[\mathbf{f}^*|\mathbf{y}, \mathbf{x}, \mathbf{x}^*, \Phi] &= K(\mathbf{x}^*, \mathbf{x}^*) - \\ &\quad [K(\mathbf{x}^*, \mathbf{x})(K(\mathbf{x}, \mathbf{x}) + \sigma^2 \mathbb{I})^{-1} \\ &\quad \times K(\mathbf{x}, \mathbf{x}^*)], \end{aligned} \quad (12)$$

where $\Phi = \{\phi, \sigma\}$.

The hyperparameters Φ are inferred as point estimates by optimising the logarithm of the marginal likelihood, with a zero mean assumption [16]:

$$\begin{aligned} \log p(\mathbf{y}|\Phi) &= \log \int p(\mathbf{y}|\mathbf{f}, \Phi) p(\mathbf{f}|\Phi) d\mathbf{f} \\ &= -\frac{1}{2} \mathbf{y}^T (K(\mathbf{x}, \mathbf{x}) + \sigma^2 \mathbb{I})^{-1} \mathbf{y} \\ &\quad - \frac{1}{2} \log |K(\mathbf{x}, \mathbf{x}) + \sigma^2 \mathbb{I}| - \frac{L}{2} \log(2\pi). \end{aligned} \quad (13)$$

4.1 Derivative Gaussian processes

Since differentiation is a linear operator, the derivative of a Gaussian process is another Gaussian process [17]. This makes it possible to include derivative observations in the GP model, or to compute predictions about derivatives. We have

$$E\left[\frac{\partial f(x)}{\partial x}\right] = \frac{\partial E[f(x)]}{\partial x}. \quad (14)$$

And likewise the covariance between partial derivative and a function value can be written as

$$K\left(\frac{\partial f(x)}{\partial x}, f(x^*)\right) = \frac{\partial}{\partial x} K(x, x^*), \quad (15)$$

and the covariance between partial derivatives follows

$$K\left(\frac{\partial f(x)}{\partial x}, \frac{\partial f(x^*)}{\partial x^*}\right) = \frac{\partial^2}{\partial x \partial x^*} K(x, x^*). \quad (16)$$

For example considering the squared exponential covariance function given in (9), we can write the covariance between partial derivative and a function value as

$$K\left(\frac{\partial f(x)}{\partial x}, f(x^*)\right) = -\frac{(x - x^*)}{l^2} K(x, x^*) \quad (17)$$

The conditional distribution of the derivative of a test function $\frac{\partial f(x^*)}{\partial x^*}$ having observed the output $(x, y) = (x, f(x))$ of a noisy function f is given as

$$P\left(\frac{\partial f(x^*)}{\partial x^*} | f(x)\right) = \mathcal{N}(\mathbf{m}, \Sigma), \quad (18)$$

where we have the mean function \mathbf{m} and covariance function Σ (obtained by using (14) and (15) and considering the prior mean of the test and training function to be zero) given as

$$\begin{aligned} \mathbf{m} &= \frac{\partial K(x^*, x)}{\partial x^*} [K(x, x) + \sigma^2 \mathbb{I}]^{-1} f(x), \\ \Sigma &= \frac{\partial^2 K(x^*, x^*)}{\partial x^* \partial x^*} \\ &\quad - \frac{\partial K(x^*, x)}{\partial x^*} [K(x, x) + \sigma^2 \mathbb{I}]^{-1} \frac{\partial K(x, x^*)}{\partial x^*}. \end{aligned} \quad (19)$$

5 ABC-SMC WITH DERIVATIVE GP

In this section we apply the machinery reviewed in the previous sections to the task of inferring parameters in differential equation models whose solution is the state trajectory $\mathbf{X}(t)$. If we assign a GP prior to the state evaluated at time points $t^L \triangleq \{t_i\}_{i=1, \dots, L}$, then the set of values of the state $\mathbf{X}(t^L)$ takes on a Gaussian **prior** distribution:

$$p(\mathbf{X}(t^L) | t^L) = \mathcal{N}(\mathbf{X}(t^L) | 0, K(t^L, t^L)). \quad (20)$$

The modelling task is to represent the experimental data as (3) $Y^d = \{\mathbf{X}(t^L) + \boldsymbol{\eta}^L\}$ where $\boldsymbol{\eta}^L$ refers to L i.i.d. samples from $\mathcal{N}(0, \sigma^2 \mathbb{I})$ and $\mathbf{X}(t)$ satisfies a differential equation. We can use GP regression to obtain the expectation and variance of the posterior (given training data Y^d at t^L) state $\mathbf{X}(t^*)$ for some test input time point t^* as in Section 4 [16]:

$$\begin{aligned} E[\mathbf{X}(t^*) | Y^d] &= K(t^*, t^L) (K(t^L, t^L) + \sigma^2 \mathbb{I})^{-1} Y^d, \\ \text{Var}[\mathbf{X}^*] &= K(t^*, t^*) - K(t^*, t^L) (K(t^L, t^L) + \sigma^2 \mathbb{I})^{-1} K(t^L, t^*). \end{aligned} \quad (21)$$

This expected posterior state variable $\mathbf{X}(t^*)$ for arbitrary choice of t^* models the smoothed evolution of the state $\widehat{\mathbf{X}}(t)$ introduced above, and where it is assumed that observational noise accounts for deviations from the smoothed time course. The smoothed state estimation enables us to compute the velocity field using the derivative GP (as in Section 4.1):

$$E\left[\frac{d}{dt} \mathbf{X}\right] = \frac{\partial K(t, t)}{\partial t} (K(t, t) + \sigma^2 \mathbb{I})^{-1} E[\mathbf{X}]. \quad (22)$$

This completes the procedure for deriving the empirical velocity field (6) $\mathbf{V}^d(t) = E[\frac{d}{dt} \mathbf{X}]$.

To apply this derivative process within the ABC framework we need to define a distance metric $\Delta(\mathbf{V}^d(t), \mathbf{f}(\widehat{\mathbf{X}}(t), \boldsymbol{\theta}))$ between the smoothed velocity field derived from the observed data, and the velocity field postulated in a differential equation model, where the expected state estimation $\widehat{\mathbf{X}}(t)$ has been substituted

for the state variable. Hence our proposed fast alternative ABC-SMC based on GP gradient distance (GP-ABC-SMC) works as follows:

- 1) Having given data Y^d as a noisy observation of the true state variable $\mathbf{X}(t)$, assign a GP prior on $\mathbf{X}(t)$ using (7) and choose a covariance function, with some unknown hyperparameters, needed to define the GP prior.
- 2) Learn the hyperparameters of the covariance function from the original noisy experimental data Y^d using maximum likelihood estimation and then run GP regression to obtain an estimation of the smoothed state evolution $\widehat{\mathbf{X}}(t) = E[\mathbf{X}]$ using (21) and the experimental time points t as both the training and test input points.
- 3) Construct the first derivative of the covariance matrix and estimate the derivative process $\mathbf{V}^d(t) = E[\frac{d}{dt} \mathbf{X} | \mathbf{X} = \widehat{\mathbf{X}}]$ using (22).
- 4) Run the ABC-SMC algorithm with a modified distance metric $\Delta(\mathbf{V}^d(t), \mathbf{f}(\widehat{\mathbf{X}}(t), \boldsymbol{\theta})) \leq \epsilon_\tau$ for tolerance schedule $\{\epsilon_\tau\}$, where at each simulation step the simulated data $Y^s = \mathbf{f}(\widehat{\mathbf{X}}(t), \boldsymbol{\theta})$ is generated by evaluating the velocity field on the right hand side of the differential equation. This yields the posterior distribution of the parameters $p(\boldsymbol{\theta} | Y^d)$ (5).

Note that no explicit solution of differential equation is required to generate the simulated data within the iterations of the GP-ABC-SMC algorithm. Also note the fact that, in order to run the GP-ABC-SMC algorithm, no knowledge of the initial condition is required.

For the sake of conformity with the ABC terminologies, we will persist in using the terms Y^d for observed data and Y^s for simulated data, as before. However, within the context of GP-ABC-SMC algorithm, observed and simulated data refer to $\mathbf{V}^d(t)$ and $\mathbf{f}(\widehat{\mathbf{X}}(t), \boldsymbol{\theta})$ respectively. That is what Y^d and Y^s will refer to for the rest of the paper.

5.1 Algorithmic settings

The success of ABC-SMC algorithm both in terms of computational complexity and quality of the solution depends on the choice of the ϵ_τ schedule and the perturbation kernel K_τ . In this section we will briefly describe how we have chosen these two algorithmic settings. A detailed discussion concerning the effects of these settings can be found in [6].

5.1.1 Tolerance schedule

Until recently, tolerance values were manually tuned in practice based on prior empirical knowledge about the model. An adaptive choice of the tolerance values has been proposed in [18] and [19]. In an adaptive schedule the value of the tolerance ϵ_τ is chosen as the α -th quantile, where $0 \leq \alpha \leq 1$ of the distances between the observed data Y^d and simulated data $Y_{\tau-1}^s$ generated at the previous algorithmic time.

5.1.2 Perturbation kernel

Perturbation kernels hold the key to the acceptance rates in ABC-SMC and the speed of the algorithm as exploited in [6]. Perturbation kernels can be broadly divided into two categories: a component-wise perturbation kernel and a multivariate perturbation kernel. In a component-wise perturbation kernel $\theta \sim \mathcal{N}(\theta, \Sigma_\tau)$ where Σ_τ is a diagonal covariance matrix whose diagonal entries $\sigma_{\tau,j}^2$, $j = 1, \dots, d$ are chosen adaptively according to the previous population labelled by $\tau - 1$ [10], [20], [6].

A component-wise perturbation kernel is, by construction, unable to generate particles with correlated components; therefore, for models with strongly correlated parameters the ABC-SMC sample generator will not be able to reflect the structure of the posterior and the acceptance rate will be low. Thus, in order to capture such correlations the particles can be perturbed according to a multivariate normal distribution with a non-diagonal covariance matrix Σ_τ that depends on the covariance of the particles as reflected in the population in the previous sequential step ($\tau - 1$) [6]. Furthermore, a multivariate perturbation kernel operating on a subset of size N' of the N particles (a local kernel) was also shown [6] to produce a noticeable improvement in the acceptance rate. In order to define this kernel we will introduce some notation. Let $Y_\tau^{s(i)}$ denote the simulated data generated from \mathcal{M}_θ with particle $\theta \leftarrow \theta_\tau^{(i)}$, $i = 1, \dots, N$ from a population of size N at algorithmic time τ . The corresponding importance weights are denoted as $w_\tau^{(i)}$. We collect all such particles (along with the weights) from algorithmic time $\tau - 1$ for which $Y_\tau^{s(i)}$ is not only within distance $\epsilon_{\tau-1}$ of the observed data Y^d but also within distance ϵ_τ of it. We denote such particles as $\tilde{\theta}_{\tau-1}^{(j)}$:

$$\left\{ \tilde{\theta}_{\tau-1}^{(j)} \right\}_{1 \leq j \leq N'} = \left\{ \theta_{\tau-1}^{(i)} \mid \Delta(Y^d, Y_{\tau-1}^{s(i)}) \leq \epsilon_\tau, 1 \leq i \leq N \right\}, \quad (23)$$

with associated normalised weights $\tilde{w}_{\tau-1}^{(j)} \triangleq (w_{\tau-1}^{(j)} / \bar{w})$, with $\bar{w} \triangleq \sum_j w_{\tau-1}^{(j)}$.

Having defined the pairs $(\tilde{\theta}_{\tau-1}^{(j)}, \tilde{w}_{\tau-1}^{(j)})$ we can now use a multivariate normal distribution $\mathcal{N}(\theta_{\tau-1}^{(i)}, \Sigma_\tau^i)$, with a local covariance Σ_τ^i (termed the optimal local covariance in [6]), to perturb a particle $\theta_{\tau-1}^{(i)}$, where local refers to particle i . This covariance is given by

$$\Sigma_\tau^i = \sum_{j=1}^{N'} \tilde{w}_{\tau-1}^{(j)} \left(\tilde{\theta}_{\tau-1}^{(j)} - \theta_{\tau-1}^{(i)} \right) \left(\tilde{\theta}_{\tau-1}^{(j)} - \theta_{\tau-1}^{(i)} \right)^T. \quad (24)$$

In the next section we will implement the GP-ABC-SMC algorithm to infer parameters of some standard non-linear differential equations through some toy examples and in that process we will compare and contrast our GP gradient based approach to that of ABC-SMC algorithm with explicit integration.

6 EVALUATION OF THE ALGORITHM ON BENCHMARK MODELS

To evaluate the GP-ABC-SMC algorithm we have chosen three benchmarking differential equations: The Lotka Volterra predator-prey model [21], the Hes1 loop model [22] and signal transduction cascade model [4],[3]. Each is a set of non-linear differential equations modelling biological systems and show non-trivial dynamical phenomenon such as limit cycle oscillations and non-stationary time evolution. For all these examples we have used the distance function given by (4) and have run the ABC-SMC algorithm with explicit integration using a component-wise univariate normal kernel (ABC-SMC-Comp) [5] as well as a multivariate normal kernel with the optimal local covariance matrix (ABC-SMC-OLCM) [6]. For our proposed GP-ABC-SMC we have also used both the aforementioned perturbation kernels. We refer these as the GP-ABC-SMC and GP-ABC-OLCM respectively. We believe a comparison between these four variants of ABC-SMC is required to capture the difference in speed of execution between the GP based ABC-SMC and the previous approaches reported in [5], [6], while comparing posterior estimates of the parameters. For all the examples presented here, we ran all these variants of ABC-SMC, including the proposed GP based ones, with $N = 100$ particles using an adaptive tolerance schedule set to the $\alpha = 0.1$ quantile of the distances in the previous populations. The ABC-SMC routines are written in MATLAB and for the GP regressions the GPML package [23] for MATLAB is used in the predator-prey and Hes1 loop example. For the signal transduction cascade model the GPML toolbox for MATLAB <https://github.com/SheffieldML/GPml> is used which has an implementation of the multi-layer perceptron (MLP) covariance kernel [3], required to handle the non-stationarity of some of the state variables. The explicit integrations are carried out using MATLAB's built in ODE and DDE solver routines.

6.1 ODE: The predator prey model

The Lotka Volterra [21] model depicts an ecological system that is used to describe the interaction between a predator and prey species. This ODE given by

$$\begin{aligned} \dot{x} &= \alpha x - xy \\ \dot{y} &= \beta xy - y, \end{aligned} \quad (25)$$

shows limit cycle behaviour and has been used for benchmarking in [5], [2]. $\theta = (\alpha, \beta)$ is the set of parameters and $\mathbf{X}(t) = (x(t), y(t))$ is the state vector comprising the concentrations of the predator and the prey species respectively. To create a realistic dataset we generated 11 uniformly spaced samples between the time interval $(0 \leq t \leq 10)$ from the model with parameters $\theta = (1, 1)$ and added random Gaussian noise with zero mean and standard deviation $\sigma = 0.5$ to each point. The initial values of the ODE for generating the synthetic data are

chosen as $\mathbf{X}(t=0) = (1.0, 0.5)$. In order to inspect the consistency of our proposed algorithm we created two more datasets obtained by adding two other realizations of the random noise to the ODE time courses. Thus we have three sets of artificial data (denoted as Dataset 1, 2 and 3), each of which has been corrupted by Gaussian noise with zero mean and standard deviation $\sigma = 0.5$ and sampled separately. Note that the GP-ABC-SMC algorithm does not require the estimation of additional nuisance parameters related to the initial values. The time evolution of the state and its derivative is predicted through the GP regression as described in Section 5. We have used the squared exponential covariance function given by (9) for the GP regression in this example.

From the synthetic data we perform the task of parameter inference using the four different variants of ABC-SMC discussed in the last section to compare their performance. Both α and β are chosen from uniform prior distributions $\mathcal{U}(-10, 10)$ in all cases. The number of algorithmic iterations, the value of S_{MC} is set to $S_{MC} = 6$ for the ABC-SMC-Comp and ABC-SMC-OLCM while it is set to $S_{MC} = 5$ for GP-ABC-SMC and GP-ABC-OLCM. The values differ because we have chosen these on the criteria of minimum number of adaptive iteration required for estimating a reliable posterior distribution. As the ABC-SMC with integration and the GP based variants operate on different spaces thus setting same values for S_{MC} does not produce comparable results. The specific values of S_{MC} for this and subsequent examples are chosen on the basis of multiple trials of all the four ABC-SMC algorithm on each of the datasets. The resulting parameter estimates are listed in Table 1 and the evaluation of the performance in Table 2 (left). We show in Table 1 the mean with 95% confidence intervals of the last population of parameters, approximating the marginal posterior, for each variants of the ABC-SMC. Note that the run-time of the GP-ABC-SMC and the GP-ABC-OLCM algorithms is the sum of the run-time of the ABC and the GP regression (including the estimation of covariance hyperparameters). The value of σ is estimated as part of the GP regression. These estimated values are $\sigma = \{0.4752, 0.8090\}$, $\sigma = \{0.6219, 0.3940\}$ and $\sigma = \{0.6432, 0.4592\}$ for the dataset 1, 2 and 3 respectively. The results show that the GP-ABC-SMC and the GP-ABC-OLCM are considerably faster than the other two variants while producing similar results in terms of the mean of the estimates. However, the variants of ABC-SMC with explicit integration produces narrower confidence intervals compared to the GP variants. Interestingly the GP variants generate comparatively fewer number of samples, as evident from Table 2 (left), indicating higher parameter sensitivities in the derivative space. Also note that the variants of ABC-SMC with a local multivariate kernel generate fewer particles compared to the ones with univariate perturbation kernels, reducing the run-times. Since dimensionality of the local covariance is very small (equal to the ODE parameter space) much less effort is required in its computation

than generating simulated data, even in the derivative space.

6.2 DDE: The Hes1 model

Our proposed algorithm is also able to estimate parameters of delay differential equations. The Hes1 model system is used in systems biology to provide a simplified account of the oscillatory behaviour of the concentrations $(\mu(t), p(t))$ of a species of mRNA and its corresponding protein. The model, introduced in [22], is described by the following delay differential equations:

$$\begin{aligned}\dot{\mu} &= \frac{1}{1 + (p(t - t_d)/p_0)^n} - \mu_m \mu \\ \dot{p} &= \mu - \mu_p p,\end{aligned}\tag{26}$$

where the parameters μ_m and μ_p are decay rates, p_0 is the repression threshold, n is the Hill coefficient and t_d is the time delay. We generated data from the above model with parameters $\mu_m = 0.03$, $\mu_p = 0.03$, $p_0 = 100$ and $t_d = 25$ and initial conditions $\mu(t_0) = 3$ and $p(t_0) = 3$ for the concentrations between the interval $(0 \leq t \leq 300)$ with uniform spacing of $\Delta t = 2$ by numerically solving the DDE. n is fixed at a value of 5 [22]. We estimated the standard deviations $\sigma_\mu = 6.0020$ and $\sigma_p = 121.7670$ of the generated data, for each of the concentrations $\mu(t)$ and $p(t)$. We then added noise, with standard deviation set to 0.1 times these estimated standard deviations σ_μ and σ_p , to the data to create the artificial datasets. As in the previous example we created three datasets in a similar fashion.

For comparison of performance of the four methods in the parameter estimation task, we keep the same algorithmic settings, as well as the same covariance function for the GP regression as in the previously example. Unlike the ODE case where our algorithm does not need to guess the initial state values, it does need a history function for $\mathbf{X}(t \leq 0)$ for DDEs in order to work. In most practical cases the initial history function is taken as a constant function. Thus in order to make our algorithm work, we shifted the first element of the estimated state evolution backward in time to create the history function. The four variants of ABC-SMC having the same settings as before, are applied to this artificial dataset. We chose uniform priors for each of the parameters: $\mu_m \sim \mathcal{U}(-2, 2)$, $\mu_p \sim \mathcal{U}(-2, 2)$, $p_0 \sim \mathcal{U}(0, 200)$ and $t_d \sim \mathcal{U}(0, 50)$. The number of iterations are chosen as $S_{MC} = 14$ while running ABC-SMC-Comp and ABC-SMC-OLCM. For GP-ABC-SMC and GP-ABC-OLCM this is chosen as $S_{MC} = 9$. As in the previous example these S_{MC} values are found through multiple trials of each of the algorithms on these datasets and inspecting the quality of the posterior estimates. The results are listed in Table 3 and Table 2 (right). In this example we see a huge speedup while using our proposed GP-ABC-SMC and GP-ABC-OLCM algorithms, demonstrating the benefits of this approach. As in the previous example we noticed higher acceptance rates (fewer generated particles) for

TABLE 1

Estimated parameters of the Lotka Volterra predator-prey model denoted by the mean of the particles of the final population and 95% confidence intervals for datasets 1-3 respectively in each row.

Parameters	True value	ABC-SMC-Comp	ABC-SMC-OLCM	GP-ABC-SMC	GP-ABC-OLCM
α	1	1.0688 ± 0.0021	1.0722 ± 0.0032	1.0373 ± 0.0080	1.0356 ± 0.0070
		1.0389 ± 0.0062	1.0305 ± 0.0054	1.0843 ± 0.0092	1.0718 ± 0.0102
		1.0293 ± 0.0157	1.0421 ± 0.0073	1.0103 ± 0.0172	1.0144 ± 0.0139
β	1	0.9698 ± 0.0025	0.9763 ± 0.0033	0.9567 ± 0.0083	0.9540 ± 0.0067
		0.9749 ± 0.0093	0.9966 ± 0.0082	0.9846 ± 0.0080	0.9760 ± 0.0083
		1.0103 ± 0.0206	0.9924 ± 0.0095	0.9887 ± 0.0154	0.9906 ± 0.0126

TABLE 2

Run-time and the ratio of total number of particles accepted to that of generated for the four ABC-SMC algorithms when applied to the three artificial datasets pertaining to the Lotka Volterra predator-prey model (left) and Hes1 model (right). The values for run-time are rounded to nearest integers.

Algorithms	Run-time (seconds)	Accept/Generate	Algorithms	Run-time (seconds)	Accept/Generate
ABC-SMC-Comp	397	700/14737	ABC-SMC-Comp	106980	1300/763045
	477	600/12294		840330	1600/8026943
	516	500/13381		201710	1600/1656197
ABC-SMC-OLCM	184	800/7846	ABC-SMC-OLCM	5496	1300/31342
	221	700/6369		8399	1500/50968
	212	600/6086		5999	1400/36519
GP-ABC-SMC	25	500/7650	GP-ABC-SMC	30	1100/31439
	26	500/7547		38	1000/38911
	26	500/7642		32	1000/36521
GP-ABC-OLCM	21	500/4655	GP-ABC-OLCM	18	1000/7387
	20	500/4316		18	1000/8183
	16	500/3193		17	1000/7969

TABLE 3

Estimated parameters of the Hes1 loop model.

Parameters	True value	ABC-SMC-Comp	ABC-SMC-OLCM	GP-ABC-SMC	GP-ABC-OLCM
μ_m	0.03	$0.0307 \pm 2.1608 \times 10^{-4}$	$0.0305 \pm 2.9381 \times 10^{-4}$	$0.0295 \pm 1.3929 \times 10^{-4}$	$0.0293 \pm 1.1127 \times 10^{-4}$
		$0.0341 \pm 6.5563 \times 10^{-5}$	$0.0342 \pm 4.3428 \times 10^{-5}$	$0.0304 \pm 2.6599 \times 10^{-4}$	$0.0302 \pm 2.3190 \times 10^{-4}$
		$0.0336 \pm 4.1649 \times 10^{-5}$	$0.0336 \pm 8.5782 \times 10^{-5}$	$0.0291 \pm 1.8721 \times 10^{-4}$	$0.0292 \pm 1.8657 \times 10^{-4}$
μ_p	0.03	$0.0294 \pm 2.0457 \times 10^{-4}$	$0.0297 \pm 2.8397 \times 10^{-4}$	$0.0300 \pm 3.8592 \times 10^{-6}$	$0.0300 \pm 2.8604 \times 10^{-6}$
		$0.0267 \pm 4.4742 \times 10^{-5}$	$0.0267 \pm 2.9662 \times 10^{-5}$	$0.0300 \pm 4.1123 \times 10^{-6}$	$0.0300 \pm 3.0506 \times 10^{-6}$
		$0.0268 \pm 2.9233 \times 10^{-5}$	$0.0268 \pm 5.9925 \times 10^{-5}$	$0.0297 \pm 3.5370 \times 10^{-6}$	$0.0297 \pm 3.6808 \times 10^{-6}$
p_0	100	99.4130 ± 0.0537	99.4518 ± 0.0697	99.5991 ± 0.2946	99.6997 ± 0.2058
		102.1872 ± 0.0362	102.2306 ± 0.0281	100.8624 ± 0.2628	100.8624 ± 0.2302
		101.2097 ± 0.0279	101.2549 ± 0.0495	100.0403 ± 0.2796	100.0593 ± 0.2459
t_d	100	25.1318 ± 0.0109	25.1580 ± 0.0151	25.0496 ± 0.1139	25.0502 ± 0.0871
		25.2317 ± 0.081	25.2428 ± 0.0056	25.9357 ± 0.2031	25.6215 ± 0.1589
		25.0730 ± 0.0055	25.0714 ± 0.0107	25.3187 ± 0.1414	25.4469 ± 0.1519

the GP variants of ABC-SMC. The noise is estimated as $\sigma_\mu = 6.8080$, $\sigma_p = 128.6910$, $\sigma_\mu = 6.9280$, $\sigma_p = 123.2920$ and $\sigma_\mu = 6.1220$, $\sigma_p = 128.1290$ for the dataset 1, 2 and 3 respectively. Furthermore, from Table 3 it is apparent that although the means of the parameters have similar estimates, their corresponding confidence intervals are different between the proposed GP variants and the original ABC-SMC algorithms. However it should be considered that for GP-ABC-SMC (with both the perturbation kernels) no knowledge of the initial history function is required. Thus in a practical setting we believe a GP based ABC-SMC algorithm is the optimal choice among these four methods for parameter estimation in DDEs.

6.3 ABC variability

Our proposed method comprises of two levels of approximation, one induced through the GP regression and the other one resulting from the approximate inference scheme. Thus in order to check the robustness of our proposed algorithm we repeated the GP-ABC-SMC and GP-ABC-OLCM parameter inference steps for 50 runs on each of the three artificial datasets for both the Lotka Volterra and Hes1 models. We used the same algorithmic settings and prior distributions as in the previous examples. Fig. 1. and Fig. 2. summarizes the distributions of the sample mean and variance (corresponding to the final particle population for each run of GP-ABC-SMC and GP-ABC-OLCM on the three artificial datasets) across all the 50 runs on the data from Lotka Volterra and


Hes1 respectively.

It is evident from Fig. 1. that the GP-ABC-OLCM algorithm produces fewer outliers compared to the GP-ABC-SMC for both the mean and variance estimates. This can be attributed to the local moves in the parameter space caused by the multivariate (OLCM) perturbation kernel. Furthermore, note that the distributions of the variances resulting from each of the algorithms are skewed in a similar manner with the outliers located at the same direction (for both GP variants) from the median. Thus these outliers represent greater variance of posterior distribution of the parameters. However, in more than 90% out of the 50 runs the moments for both the parameters lie within the inter-quantile range.


In case of the Hes1 model it is apparent from Fig. 2. that the distributions are less variable across multiple runs and variants of the algorithms. Moreover, in this case we notice that the distribution of the variances have very few outliers indicating greater accordance among the posteriors learnt after each run of the algorithms. However, it should be noted that an adaptive tolerance schedule results in different (marginally) tolerance values for each new run of the ABC-SMC. Thus some amount of variability in the moments corresponding to different runs is attributed to the differing tolerances. Fig. 3. and Fig. 4. show the learnt state trajectories of the Lotka Volterra and Hes1 model compared against the true state trajectories for each of the datasets. The true trajectories correspond to the true parameters and the reconstructed trajectories are generated by solving the Lotka Volterra (25) and Hes1 (26) model equations. While solving (numerically integrating) these differential equations the parameters are taken as the median of the parameters learnt by the GP-ABC-SMC algorithm considering all the 50 runs. The median value is considered here to reflect the effect of variability (in parameter learning by the GP-ABC-SMC) in reconstructing the dynamics of the considered models.

6.4 Signal transduction cascade

We have, so far, used the benchmarking examples to compare our proposed GP-based ABC-SMC approach to others of that ilk that exist in the literature. In this example we will compare the parameter estimation results for the proposed GP based ABC-SMC with other (methods not falling under ABC) recent GP based approximate inference methods for parameter estimation in ODEs. For this purpose we have chosen the signal transduction cascade model [4]. Using this model, a comparison between the competing GP based approaches were reported in [3]. Thus evaluating the proposed GP based ABC-SMC algorithm on this model (with identical settings to those in [3]) will enable us to draw comparisons with these other methods. This model is described by a 5-dimensional coupled ODEs given by



(a) Distribution of the mean for Lotka Volterra across 50 runs of GP-ABC-SMC and GP-ABC-OLCM.

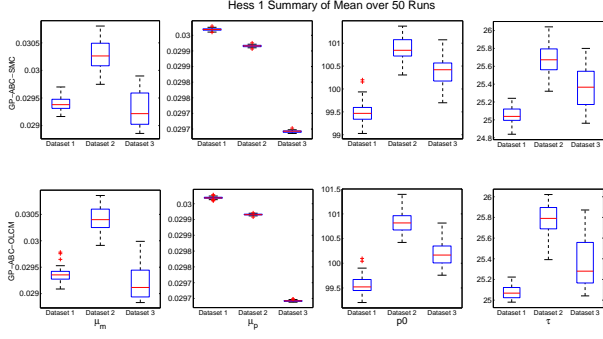


(b) Distribution of the variance for Lotka Volterra across 50 runs of GP-ABC-SMC and GP-ABC-OLCM.

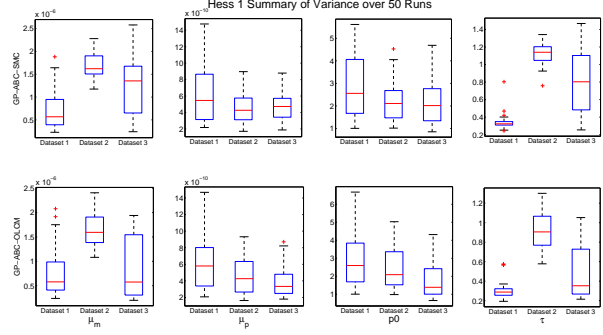
Fig. 1. The boxplots represent the distributions of the mean and variances (across 50 runs) of the final population representing the marginal approximate posterior parameter distributions learnt by the GP-ABC-SMC and the GP-ABC-OLCM from the three artificial datasets of Lotka Volterra model.

$$\begin{aligned}
 \frac{d[S]}{dt} &= -k_1[S] - k_2[S][R] + k_3[RS] \\
 \frac{d[S_d]}{dt} &= k_1[S] \\
 \frac{d[R]}{dt} &= -k_2[S][R] + k_3[RS] + \frac{V[R_{pp}]}{K_m + [R_{pp}]} \\
 \frac{d[RS]}{dt} &= k_2[S][R] - k_3[RS] - k_4[RS] \\
 \frac{d[R_{pp}]}{dt} &= k_4[RS] - \frac{V[R_{pp}]}{K_m + [R_{pp}]}
 \end{aligned} \tag{27}$$

where $\theta = (k_1, k_2, k_3, k_4, V, k_m)$ are the parameters of this model and $\mathbf{X}(t) = ([S], [S_d], [R], [RS], [R_{pp}])$ are the concentrations of the state variables. Following [3] we generated data from the model between the time interval $(0 \leq t \leq 100)$ with parameters $\theta = (0.07, 0.6, 0.05, 0.3, 0.017, 0.3)$ and initial values of the state variable $[S] = 1, [S_d] = 0, [R] = 1, [RS] = 0, [R_{pp}] = 0$. We then sampled the data at time $t^L = \{0, 1, 2, 4, 5, 7, 10, 15, 20, 30, 40, 50, 60, 80, 100\}$ and added



(a) Distribution of the mean for Hes1 across 50 runs of GP-ABC-SMC and GP-ABC-OLCM.



(b) Distribution of the variance for Hes1 across 50 runs of GP-ABC-SMC and GP-ABC-OLCM.

Fig. 2. Distributions of the mean and variances learnt by the GP-ABC-SMC and the GP-ABC-OLCM from the three artificial datasets of Hes1 model.

random noise with standard deviation $\sigma_{[S]}$, $\sigma_{[S_d]}$, $\sigma_{[R]}$, $\sigma_{[RS]}$, $\sigma_{[R_{pp}]}$ set to 0.1 for generating the synthetic data. For inferring parameters in this example we apply the GP-ABC-OLCM algorithm from our study with multiple runs, where we found this algorithm to provide a stable and fast inference mechanism. The non-stationarity in the time evolution of the state variables is captured by the MLP covariance function given by

$$k(t, t') = \sigma_{\text{kernel}}^2 \times \frac{2}{\pi} \arcsin \left(\frac{\sigma_w^2 t^\top t' + \sigma_b^2}{\sqrt{\sigma_w^2 t^\top t + \sigma_b^2 + 1} \sqrt{\sigma_w^2 t'^\top t' + \sigma_b^2 + 1}} \right), \quad (28)$$


where the kernel variance σ_{kernel}^2 , the neural network weight variance σ_w^2 , and the bias variance σ_b^2 are the hyperparameters of the covariance function. The derivative of this kernel with respect to the input time t is given by

$$\frac{\partial k(t, t')}{\partial t} = \frac{\sigma_{\text{kernel}}^2}{\sqrt{1 - Z^2}} \frac{\partial Z}{\partial t}, \quad (29)$$


where

$$Z = \frac{\sigma_w^2 t^\top t' + \sigma_b^2}{Z_{\text{norm}}} \quad (30)$$

with $Z_{\text{norm}} = \sqrt{\sigma_w^2 t^\top t + \sigma_b^2 + 1} \sqrt{\sigma_w^2 t'^\top t' + \sigma_b^2 + 1}$. All the other algorithmic settings were kept the same. The prior distributions are chosen as $k_1 \sim \mathcal{U}(0.05, 0.09)$,



(a) Trajectories of $x(t)$.



(b) Trajectories of $y(t)$.


Fig. 3. Reconstructed and true state trajectories of the Lotka Volterra model. The results corresponding to the three datasets (1, 2, 3) are shown using blue, red and magenta colours respectively. Reconstructed trajectories are represented as curves and observations as stars. The ground truth is the black (dashed and circled) curve.

$k_2 \sim \mathcal{U}(0.4, 0.8)$, $k_3 \sim \mathcal{U}(0.03, 0.07)$, $k_4 \sim \mathcal{U}(0.1, 0.5)$, $V \sim \mathcal{U}(0.015, 0.0195)$ and $k_m \sim \mathcal{U}(0.1, 0.5)$. In this example S_{MC} is set to 3.


The resulting parameter estimates are furnished in Table 4 along with the parameter estimates obtained from other GP based algorithms run on the same model. These algorithms are the **GP-ODE** method proposed in [3], the adaptive gradient matching (**AGM**) proposed in [2] and the gradient matching (**GM**) proposed in [1]. The posterior distributions of the parameters as learnt by the GP-ABC-OLCM is shown in Fig. 6. We have compared the true state trajectories with the reconstructed trajectories in Fig. 5. We generated the reconstructed trajectories by solving (27) using the mean of the final population of GP-ABC-OLCM, representing the marginal posterior densities of the parameters. The estimated values of the standard deviations are $\sigma_{[S]} = 0.0964$, $\sigma_{[S_d]} = 0.0818$, $\sigma_{[R]} = 0.0707$, $\sigma_{[RS]} = 0.0591$ and $\sigma_{[R_{pp}]} = 0.0754$. We avoided the comparison of run-time or acceptance rates as the GP-ABC-OLCM and other GP based algorithms depend on completely different approximate inference scheme. However, GP-ABC-OLCM is significantly faster than the other approaches. The GP-ABC-OLCM finishes

TABLE 4
Estimated parameters of the signal transduction cascade by all the GP based approaches including the GP-ABC-OLCM. The estimates for **GP-ODE**, **AGM** and **GM** are taken from [3].

Parameters	True value	GP-ABC-OLCM	GP-ODE	AGM	GM
k_1	0.070	0.0708 ± 0.0086	0.0747 ± 0.0130	0.0771 ± 0.0130	0.0762 ± 0.0130
k_2	0.6	0.5806 ± 0.0706	0.6230 ± 0.1246	0.5460 ± 0.1259	0.5632 ± 0.1256
k_3	0.05	0.0480 ± 0.0074	0.0530 ± 0.0135	0.0593 ± 0.0111	0.0594 ± 0.0115
k_4	0.3	0.3439 ± 0.0659	0.2960 ± 0.0281	0.3750 ± 0.0999	0.3754 ± 0.1051
V	0.017	0.0170 ± 0.0009	0.0177 ± 0.0014	0.0172 ± 0.0015	0.0173 ± 0.0014
k_m	0.3	0.3110 ± 0.0774	0.4220 ± 0.0690	0.4090 ± 0.0911	0.4186 ± 0.0953



(a) Trajectories of $\mu(t)$.



(b) Trajectories of $p(t)$.

Fig. 4. Reconstructed and true state trajectories of the Hes1 model.

the estimation in around 20 seconds while the other methods were run for 30 minutes to obtain a properly mixed Markov chain. The ratio of the last two parameters V/k_m [2] is a crucial quantity that determines the reconstruction accuracy. GP-ABC-OLCM is able to infer this quantity with the best (based on the estimated posterior means of V and k_m) accuracy among all the GP based algorithms.

7 CONCLUSION

In this paper we have proposed a method that significantly speeds up the task of parameter inference in comparison with state of the art methods that use ABC and SMC based approaches when applied to several benchmark dynamical system models that are described

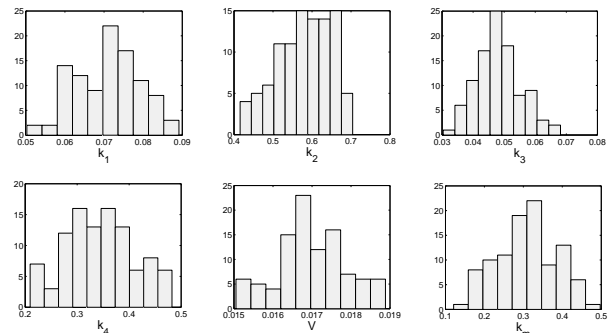


Fig. 6. Posterior distributions of the parameters of the signal transduction cascade model learnt by the proposed GP-ABC-OLCM algorithm.

by ordinary and delay differential equations. We achieve this speed-up by circumventing the need to numerically integrate the differential equations, a task that is repeatedly required in other ABC methods to generate samples from candidate models for comparison with the observed data. The key idea behind our method lies in building on [1], [12] and [3] to work directly with the vector field of the dynamical system, which we model using Gaussian process regression, and thus create a distance function in derivative space for use in the ABC-SMC algorithm as proposed in [5]. Thus we proposed a modified ABC-SMC algorithm for parameter estimation in ODEs or DDEs.

We benchmarked the benefits of this approach by evaluating our proposed approach on toy problems where we observed a significant speed-up of the parameter estimation process. We also compared our proposed approach with other GP based methods proposed in recent literature and found that our proposed GP-ABC-SMC (with the local multivariate perturbation kernel) performs significantly faster to obtain similar estimates. Furthermore, improvements of ABC-SMC through perturbation kernels as proposed in [6] and has been integrated with our approach to obtain enhanced performance. Thus, our fast approximate inference process can accommodate the useful features of ABC-SMC (as shown in [5]) such as model selection and sensitivity analysis. Our proposed approach is only limited by the ability of the Gaussian process regression in smoothing the observed time series data while retaining the essential

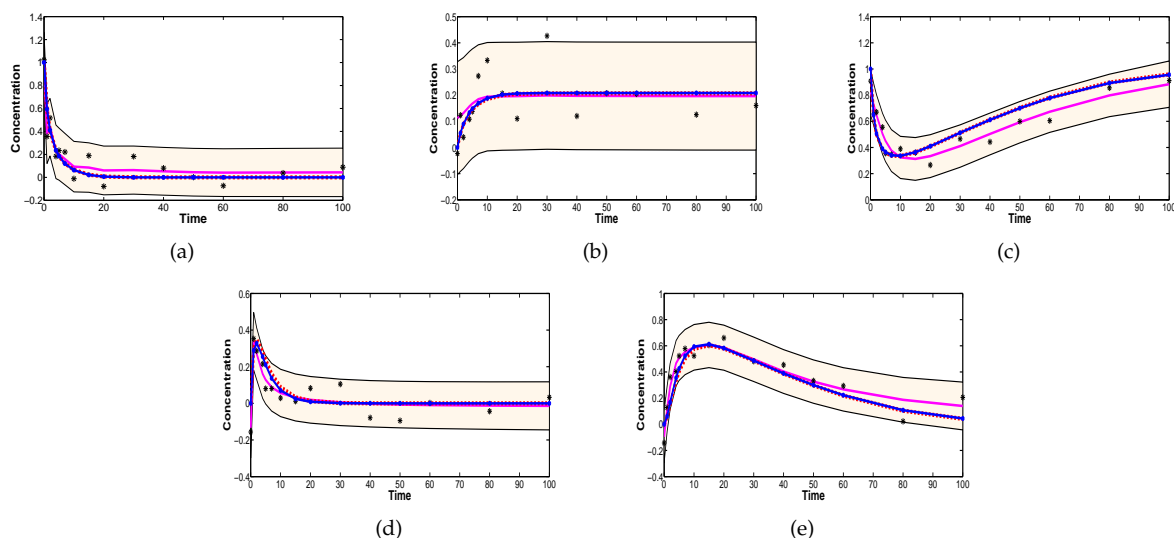


Fig. 5. Results of GP-ABC-OLCM for the signal transduction cascade ($([S], [S_d], [R], [RS], [R_{pp}]$ in (a),(b),(c),(d) and (e) respectively). In all the plots observations are the black stars, the true state trajectory is the red (dashed) curve and the blue curve shows the reconstructed trajectory. We have also plotted the GP mean function as the magenta curve and GP variance as the shaded area. The reconstructed trajectory is generated by numerically integrating (27) with the parameters set to the mean of the posterior distribution estimated by the GP-ABC-OLCM algorithm.

characteristics that are meant to be captured by the dynamical system model. Thus in those cases where smoothing the experimental data by GP regression introduces artefacts, the GP-ABC-SMC algorithm would produce poor parameter estimates. As future work we wish to extend the GP-ABC-SMC algorithm for stochastic differential equation by relating the GP regression technique with drift estimation technique in [24]. We will also include models with hidden variables, as the smoothing procedure on observed data can no longer provide complete information.

ACKNOWLEDGMENTS

The work reported in this paper was supported by project PLants Employed As Sensor Devices (PLEASED), EC grant agreement number 296582.

REFERENCES

- [1] B. Calderhead, M. Girolami, and N. D. Lawrence, "Accelerating Bayesian Inference over Nonlinear Differential Equations with Gaussian Processes," in *Proc. Advances in Neural Information Processing Systems 21*, 2008, pp. 217–224.
- [2] F. Dondelinger, M. Filippone, S. Rogers, and D. Husmeier, "ODE parameter inference using adaptive gradient matching with Gaussian processes," in *Proc. 16th Int'l Conf. on Artificial Intelligence and Statistics*, 2013, pp. 216–228.
- [3] Y. Wang and D. Barber, "Gaussian Processes for Bayesian Estimation in Ordinary Differential Equations," in *Proc. 31st Int'l. Conf. Machine Learning*, 2014, pp. 1485–1493.
- [4] V. Vyshemirsky and M. A. Girolami, "Bayesian ranking of biochemical system models," *Bioinformatics*, vol. 24, no. 6, pp. 833–839, 2008.
- [5] T. Toni, D. Welch, N. Strelkowa, A. Ipsen, and M. P. Stumpf, "Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems," *Journal of the Royal Society Interface*, vol. 6, no. 31, pp. 187–202, Feb. 2009.
- [6] S. Filippi, C. P. Barnes, J. Cornebise, and M. P. Stumpf, "On optimality of kernels for approximate bayesian computation using sequential monte carlo," *Statistical applications in genetics and molecular biology*, vol. 12, no. 1, pp. 87–107, 2013.
- [7] S. A. Sisson, Y. Fan, and M. M. Tanaka, "Sequential monte carlo without likelihoods," *Proceedings of the National Academy of Sciences*, vol. 104, no. 6, pp. 1760–1765, 2007.
- [8] P. Del Moral, A. Doucet, and A. Jasra, "Sequential Monte Carlo samplers," *Journal of the Royal Statistical Society: Series B*, vol. 68, no. 3, pp. 411–436, Jun. 2006.
- [9] S. A. Sisson, Y. Fan, and M. M. Tanaka, "Correction for sisson sequential monte carlo without likelihoods," *Proceedings of the National Academy of Sciences*, vol. 106, no. 39, p. 16889, 2009.
- [10] M. A. Beaumont, J. M. Cornuet, J. M. Marin, and C. P. Robert, "Adaptive approximate Bayesian computation," *Biometrika*, vol. 96, no. 4, pp. 983–990, 2009.
- [11] J. Varah, "A spline least squares method for numerical parameter estimation in differential equations," *SIAM Journal on Scientific and Statistical Computing*, vol. 3, no. 1, pp. 28–46, 1982.
- [12] J. O. Ramsay, G. Hooker, D. Campbell, and J. Cao, "Parameter estimation for differential equations: a generalized smoothing approach," *Journal of the Royal Statistical Society: Series B*, vol. 69, no. 5, pp. 741–796, Nov. 2007.
- [13] A. O'Hagan and J. Kingman, "Curve fitting and optimal design for prediction," *Journal of the Royal Statistical Society: Series B*, vol. 40, pp. 1–42, 1978.
- [14] D. J. MacKay, "Introduction to gaussian processes," *NATO ASI Series F Computer and Systems Sciences*, vol. 168, pp. 133–166, 1998.
- [15] R. M. Neal, "Regression and Classification Using Gaussian Process Priors," *Bayesian Statistics*, vol. 6, pp. 475–501, 1998.
- [16] C. E. Rasmussen and C. K. I. Williams, *Gaussian processes for machine learning*. MIT Press, 2006.
- [17] E. Solak, R. Murray-Smith, W. Leithead, C. Rasmussen, and D. Leith, "Derivative observations in Gaussian process models of dynamic systems," in *Proc. Advances in Neural Information Processing Systems 15*, 2002, pp. 1033–1040.
- [18] P. Del Moral, A. Doucet, and A. Jasra, "An adaptive sequential Monte Carlo method for approximate Bayesian computation," *Statistics and Computing*, vol. 22, no. 5, pp. 1009–1020, 2012.
- [19] C. C. Drovandi and A. N. Pettitt, "Estimation of parameters for macroparasite population evolution using approximate bayesian computation," *Biometrics*, vol. 67, no. 1, pp. 225–33, Mar. 2011.
- [20] X. Didelot, R. G. Everitt, A. M. Johansen, and D. J. Lawson,

- "Likelihood-free estimation of model evidence," *Bayesian analysis*, vol. 6, no. 1, pp. 49–76, Mar. 2011.
- [21] J. Murray, *Mathematical Biology: I. An Introduction*. Springer, 2002.
 - [22] N. Monk, "Oscillatory expression of Hes1, p53, and NF- κ B driven by transcriptional time delays," *Current Biology*, vol. 13, pp. 1409–1413, 2003.
 - [23] C. E. Rasmussen and H. Nickisch, "Gaussian processes for machine learning (gpml) toolbox," *The Journal of Machine Learning Research*, vol. 9999, pp. 3011–3015, 2010.
 - [24] A. Rutter, P. Batz, and M. Opper, "Approximate Gaussian process inference for the drift of stochastic differential equations," in *Proc. Advances in Neural Information Processing Systems 26*, 2013, pp. 2040–2048.